What is claimed is:

1   1. A method of accessing a group in a clustered computer system, wherein the

2   clustered computer system includes a plurality of nodes, and wherein the group

3   includes a plurality of members resident respectively on the plurality of nodes, the

4   method comprising:

5   (a) receiving an access request on a first node in the plurality of nodes,

6   wherein the access request identifies a cluster-private group name associated

7   with the group; and

8   (b) processing the access request on the first node to initiate a group

9   operation on at least a subset of the plurality of nodes that map to the cluster-

10   private group name.

1   2. The method of claim 1, further comprising generating the access request

2   with a user job resident on the first node.

1   3. The method of claim 2, further comprising forwarding the access request to

2   a clustering infrastructure resident in the first node via a call from the user job.

1   4. The method of claim 1, further comprising:

2   (a) generating the access request with a user job resident on a second

3   node in the plurality of nodes; and

4   (b) processing the access request with a proxy job resident on the

5   second node by communicating the access request to the first node.

1   5. The method of claim 4, wherein the proxy job is a member of a cluster

2   control group, the method further comprising:

3   (a) forwarding the access request from the user job to the proxy job;

4   and

5   (b) forwarding the access request from the proxy job to a clustering

6   infrastructure resident in the second node via a call from the proxy job.

1   6. The method of claim 1, further comprising retrieving the cluster-private

2   group name with a user job by accessing a cluster-private data structure.

1   7. The method of claim 6, wherein the cluster-private data structure is resident

2   on the same node as the user job.

1   8. The method of claim 7, wherein the cluster-private data structure is

2   accessible only from the node upon which the cluster-private data structure is resident.

1   9. The method of claim 8, wherein the cluster-private data structure is

2   accessible only by jobs that are resident on the node upon which the cluster-private

3   data structure is resident.

1   10. The method of claim 1, wherein initiating the group operation comprises

2   distributing messages to a plurality of group members resident on the nodes that map

3   to the cluster-private group name.

1   11. The method of claim 10, wherein initiating the group operation further

2   comprises accessing a group address data structure to retrieve a plurality of network

3   addresses associated with the cluster-private group name, wherein distributing

4   messages to the plurality of group members includes sending a message to each of the

5   plurality of network addresses.

1   12. The method of claim 1, wherein initiating the group operation is

2   performed by a clustering infrastructure resident on the first node.

1   13. The method of claim 12, wherein initiating the group operation includes

2   retrieving with the clustering infrastructure a plurality of addresses that are mapped to

3   the cluster-private group name in a data structure that is local to the clustering

4   infrastructure.

1   14. The method of claim 1, wherein initiating the group operation includes

2   locally resolving on the first node a mapping between the cluster-private group name

3   and a plurality of addresses associated with at least the subset of the plurality of

4   nodes.

1   15. An apparatus, comprising:

2   (a) a memory accessible by a first node among a plurality of nodes in a

3   clustered computer system; and

4   (b) a program resident in the memory and executed by the first node,

5   the program configured to access a group that includes a plurality of members

6   resident respectively on the plurality of nodes by receiving an access request

7   that identifies a cluster-private group name associated with the group, and

8   processing the access request to initiate a group operation on at least a subset

9   of the plurality of nodes that map to the cluster-private group name.

1   16. The apparatus of claim 15, further comprising a user job configured to

2   generate the access request.

1   17. The apparatus of claim 16, wherein the program comprises a clustering

2   infrastructure resident on the first node.

1   18. The apparatus of claim 17, further comprising a proxy job configured to

2   forward the access request from the user job to the clustering infrastructure.

1   19. The apparatus of claim 15, further comprising:

2   (a) a cluster-private data structure configured to store the cluster-

3   private group name; and

4   (b) a user job configured to access the cluster-private data structure to

5   retrieve the cluster-private group name and generate the access request

6   therefrom.

1   20. The apparatus of claim 19, wherein the cluster-private data structure is

2   resident on the same node as the user job.

1   21. The apparatus of claim 20, wherein the cluster-private data structure is

2   accessible only from the node upon which the cluster-private data structure is resident.

- 21 -

1   22. The apparatus of claim 15, further comprising a group address data

2   structure configured to store a plurality of network addresses associated with the

3   cluster-private group name, wherein the program is configured to initiate the group

4   operation by accessing the group address data structure to retrieve the plurality of

5   network addresses and sending a message to each of the plurality of network

6   addresses.

1   23. The apparatus of claim 22, wherein the program comprises a clustering

2   infrastructure, and wherein the group address data structure is local to the clustering

3   infrastructure.

1   24. The method of claim 15, wherein the program is further configured to

2   process the access request by locally resolving on the first node a mapping between

3   the cluster-private group name and a plurality of addresses associated with at least the

4   subset of the plurality of nodes.

IBM ROC920000273US1
WH&E IBM/177

1   25. A clustered computer system, comprising:

2   (a) a plurality of nodes coupled to one another over a network;

3   (b) a group including a plurality of members resident respectively on

4 the plurality of nodes; and

5   (c) a program resident in a first node among the plurality of nodes and

6 configured to access the group by receiving an access request that identifies a

7 cluster-private group name associated with the group, and processing the

8 access request to initiate a group operation on at least a subset of the plurality

9 of nodes that map to the cluster-private group name.

1   26. A program product, comprising:

2   (a) a program resident in the memory and executed by a first node

3   among a plurality of nodes in a clustered computer system, the program

4   configured to access a group that includes a plurality of members resident

5   respectively on the plurality of nodes by receiving an access request that

6   identifies a cluster-private group name associated with the group, and

7   processing the access request to initiate a group operation on at least a subset

8   of the plurality of nodes that map to the cluster-private group name; and

9   (b) a signal bearing medium bearing the program.

1   27. The program product of claim 26, wherein the signal bearing medium

2 includes at least one of a transmission medium and a recordable medium.